

Improved Bounds for 3SUM, K-SUM, and Linear Degeneracy *

Omer Gold[†]Micha Sharir[‡]

December 17, 2015

Abstract

Given a set of n real numbers, the 3SUM problem is to decide whether there are three of them that sum to zero. Until a recent breakthrough by Grønlund and Pettie [FOCS'14], a simple $\Theta(n^2)$ -time deterministic algorithm for this problem was conjectured to be optimal. Over the years many algorithmic problems have been shown to be reducible from the 3SUM problem or its variants, including the more generalized forms of the problem, such as k -SUM and k -variate linear degeneracy testing¹. The conjectured hardness of these problems have become extremely popular for basing conditional lower bounds for numerous algorithmic problems in P. Thus, a better understanding of the complexity of the 3SUM problem and its variants, might shed more light on the complexity of a wide class of problems in P.

In this paper we show the following:

- A *deterministic* algorithm for 3SUM that runs in $O(n^2 \log \log n / \log n)$ time.
- The randomized decision tree complexity of 3SUM is $O(n^{3/2})$.
- The randomized decision tree complexity of k -variate linear degeneracy testing (k-LDT) is $O(n^{k/2})$, for any odd $k \geq 3$.

These bounds improve the ones obtained by Grønlund and Pettie, giving a faster deterministic algorithm and new randomized decision tree bounds for this archetypal algorithmic problems.

1 Introduction

The general 3SUM problem is formally defined as

3SUM: Given a finite set $A \subset \mathbb{R}$, determine whether there exist $a, b, c \in A$ such that $a + b + c = 0$.

An equivalent variant is that the input consists of three finite sets $A, B, C \subset \mathbb{R}$ of the same size, and the goal is to determine whether there are elements $a \in A, b \in B, c \in C$ such that $a + b + c = 0$. When the sets A, B, C are not of the same size, the problem is named unbalanced 3SUM.

The 3SUM problem and its variants are among the most fundamental problems in algorithm design. Although the 3SUM problem itself does not seem to have many compelling practical implications, it has been of wide interest due to numerous problems that can be reduced from it. The notion of 3SUM-Hardness is often used to describe such problems, namely, problems that

*Work on this paper has been supported by Grant 892/13 from the Israel Science Foundation, by Grant 2012/229 from the U.S.-Israeli Binational Science Foundation, by the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University.

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; omergold@post.tau.ac.il

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; michas@post.tau.ac.il

¹ Given a linear function $f(x_1, \dots, x_k) = \alpha_0 + \sum_{1 \leq i \leq k} \alpha_i x_i$ and a finite set $A \subset \mathbb{R}$, the k -variate linear degeneracy testing problem (k-LDT) is to decide whether $0 \in f(A^k)$. When $f(x_1, \dots, x_k) = \sum_{i=1}^k x_i$ the problem is called k -SUM.

are at least as hard as 3SUM. Thus, lower bounds on 3SUM imply lower bounds on dozens of other problems. Among them are fundamental problems in computational geometry [4, 9, 28, 40], dynamic graph algorithms [3, 35, 37], triangle enumeration [2, 35], and pattern matching [6, 14, 21, 35, 41].

By the time hierarchy theorem [30], there are problems in P with complexity $\Omega(n^k)$ for every fixed k . However, given a problem in P, proving an $\Omega(n^k)$ unconditional lower bound, for any $k > 1$, seems far beyond the state of the art in computational complexity theory. This has led researchers to settle on conditional lower bounds, based on the conjectured hardness of certain archetypal problems, such as 3SUM, (min, +)-matrix multiplication, and CNF-SAT. See [1–3, 6, 7, 9, 12, 20, 28, 34, 35, 37–39, 41, 42] for many examples of such conditional lower bounds.

In the last decades, starting with a study of Gajentaan and Overmars [28], it was conjectured that any algorithm for 3SUM requires $\Omega(n^2)$ time. However, a recent breakthrough by Grønlund and Pettie [29] showed that 3SUM can be solved in subquadratic time. Specifically, they gave a *deterministic* algorithm that runs in $O(n^2(\log \log n / \log n)^{2/3})$ time, and a randomized algorithm that runs in $O(n^2(\log \log n)^2 / \log n)$ expected time and with high probability. Furthermore, they showed that there is a 4-linear decision tree for 3SUM with depth $O(n^{3/2} \sqrt{\log n})$. These results raised serious doubts on the optimality of many algorithms for 3SUM-Hard problems. For example, the following problems are known to be 3SUM-Hard.

1. Given an n -point set in \mathbb{R}^2 , determine whether it contains three collinear points (Gajentaan and Overmars [28]).
2. Given n triangles in \mathbb{R}^2 , determine whether their union contains a hole, or compute the area of their union [28].
3. Given two n -point sets $X, Y \subset \mathbb{R}$, each of size n , determine whether all elements in $X + Y = \{x + y \mid x \in X, y \in Y\}$ are distinct (Barequet and Har-Peled [9]).
4. Given two n -edge convex polygons, determine whether one can be placed inside the other via translation and rotation [9].

Problems 1 and 2 are solvable in $O(n^2)$ time (see [28]). Problems 3 and 4 are solvable in $O(n^2 \log n)$ time (see [9]). In face of the new 3SUM result of Grønlund and Pettie [29], it is natural to ask whether these bounds are optimal. However, no better bounds are currently known (in spite of the improvement in [29]). Problem 3 (or its stronger variant of sorting $X + Y$) has special importance, as it is used for basing the conditional lower bounds for the problems in [9] and in [31]; these problems are therefore also classified as “(Sorting $X + Y$)-Hard”. The decision tree complexity of Problem 3 was shown to be $O(n^2)$ by Fredman [25]. It is a prominent long-standing open problem whether Problem 3 can be solved in $o(n^2 \log n)$ time (see [22]).

In view of the results in [29], the 3SUM conjecture has been replaced by a relaxed, modern variant, asserting that 3SUM cannot be solved in *strongly subquadratic time* (even in expectation), i.e., in $O(n^{2-\epsilon})$ time, for any $\epsilon > 0$. This conjecture is widely accepted and believed by the computer science community, and so is its implications for deriving lower bounds for other problems. Abboud and Williams [2] claim, based on the collective computer science community efforts, that lower bounds that are based on the relaxed 3SUM conjecture should be at least as believable as any other known conditional lower bounds for a problem in P.

This relaxed conjecture is often applied to a more restricted variant, Integer3SUM, which is defined as follows.

Integer3SUM: Given a finite set $A \subseteq \{-U, \dots, U\} \subset \mathbb{Z}$, determine whether there exist $a, b, c \in A$ such that $a + b + c = 0$.

Based on the conjecture that **Integer3SUM** requires $\Omega(n^{2-o(1)})$ time, Pătraşcu [37] proved lower bounds on triangle enumeration and on various problems in dynamic data structures. Recently, the reduction techniques of Pătraşcu were extended by Kopelowitz, Pettie, and Porat [35], implying improved lower bounds for this kind of problems. Examples for lower bounds based on this conjecture include the following:

- Given an undirected m -edge graph, enumerating up to m triangles (3-cycles) requires at least $\Omega(m^{4/3-o(1)})$ time (Pătraşcu [37]).²
- Given a sequence of m updates to a directed graph (edge insertions and deletions) and two specified vertices s, t , determining whether t is reachable from s after each update, requires at least $\Omega(m^{4/3-o(1)})$ time (Abboud and Williams [3]).
- Given an edge-weighted undirected graph, deciding whether it contains a zero-weight triangle, requires at least $\Omega(n^{3-o(1)})$ time (Williams and Williams [43]).

The **Integer3SUM** problem is clearly not harder than **3SUM**; however, any other relationship between them is unknown. Unlike **3SUM**, **Integer3SUM** can be solved using fast Fourier transform in $O(n + U \log U)$ time, which is subquadratic even for a rather large universe size U .³ Baran, Demaine, and Pătraşcu [8] showed that using randomized universe reductions, word packing, and table lookups, **Integer3SUM** can be solved in $O(n^2(\log \log n / \log n)^2)$ *expected* time and with high probability, on the word RAM, where $U = 2^w$ and $w > \log n$ is the machine word size. Recently, Chan and Lewenstein [17] showed, based on results from additive combinatorics, strongly subquadratic time algorithms for special restricted cases of **Integer3SUM**.

The **3SUM** problem was also extensively studied in its generalized form, k -variate linear degeneracy testing (**k-LDT**), formally defined as

k -LDT and k -SUM: Given a k -variate linear function $\phi(x_1, \dots, x_k) = \alpha_0 + \sum_{i=1}^k \alpha_i x_i$, where $\alpha_0, \dots, \alpha_k \in \mathbb{R}$, and a finite set $A \subset \mathbb{R}$, determine whether there exists $(x_1, \dots, x_k) \in A^k$ such that $\phi(x_1, \dots, x_k) = 0$. When ϕ is $\sum_{i=1}^k x_i$ the problem is called k -SUM.

There are simple algorithms that solve k -LDT in time $O(n^{(k+1)/2})$ when k is odd, or $O(n^{k/2} \log n)$ when k is even; see [5]. These algorithms are based on straightforward reductions to a **2SUM** problem or to an unbalanced **3SUM** problem, depending on whether k is even or odd, respectively. These are currently the best known time upper bounds for solving k -LDT. Erickson [24] showed that, for an even k , there is a k -linear decision tree with depth $O(n^{k/2})$, removing an $O(\log n)$ factor when comparing to the uniform model. The above bounds match with the seminal lower bound results of Erickson [24], and Ailon and Chazelle [5], who showed that any k -linear decision tree for solving k -LDT must have depth $\Omega(n^{k/2})$ when k is even and $\Omega(n^{(k+1)/2})$ when k is odd. In particular, any 3-linear decision tree for **3SUM** has depth $\Omega(n^2)$. Grønlund and Pettie [29] showed that using more variables per comparison leads to a dramatic improvement in the depth of the tree, which significantly beats the above lower bounds. Specifically, as will be reviewed below, they showed that there is a 4-linear decision tree for **3SUM** with depth $O(n^{3/2} \sqrt{\log n})$, and by the reduction from k -LDT to unbalanced **3SUM**, they concluded that there is a $(2k-2)$ -linear decision tree for k -LDT with depth $O(n^{k/2} \sqrt{\log n})$, for any odd $k \geq 3$.

Apart from the many lower bounds obtained from the conjectured hardness of **3SUM** and its variants, in recent years, many lower bounds were obtained from two other plausible conjectures. The first is that computing the $(\min, +)$ -product of two $n \times n$ matrices takes $\Omega(n^{3-o(1)})$ time

²By Björklund et al. [10], the exponent $4/3$ is optimal if the matrix multiplication exponent ω is 2.

³Erickson [24] credits R. Seidel with this **3SUM** algorithm (taken from [29]).

3SUM		
Authors	Bound	Model of Computation
standard	$O(n^2)$	Deterministic
Grønlund & Pettie 2014	$O(n^{3/2}\sqrt{\log n})$	Decision Tree
	$O(n^2(\log \log n)^2/\log n)$	Randomized
	$O(n^2(\log \log n/\log n)^{2/3})$	Deterministic
new	$O(n^{3/2})$	Randomized Decision Tree
	$O(n^2 \log \log n/\log n)$	Deterministic

Integer3SUM		
Authors	Bound	Model of Computation
standard	$O(n^2)$	Deterministic
Seidel 1997	$O(n + U \log U)$	Deterministic
Baran, Demaine, and Pătraşcu 2005	$O(n^2(\log \log n/\log n)^2)$	Randomized
	$O(n^2(\log w)^2/w)$	Randomized
Grønlund & Pettie 2014	$O(n^2(\log \log n/\log n)^{2/3})$	Deterministic
new	$O(n^2 \log \log n/\log n)$	Deterministic

k-LDT		
Authors	Bound	Model of Computation
standard	$O(n^{k/2} \log n)$, for even k .	Deterministic
	$O(n^{(k+1)/2})$, for odd k .	Deterministic
Erickson 1999	$O(n^{k/2})$, for even k .	Decision Tree
Grønlund & Pettie 2014	$O(n^{k/2}\sqrt{\log n})$, for odd k .	Decision Tree
new	$O(n^{k/2})$, for odd k .	Randomized Decision Tree

Figure 1.1: A summary of the results. In **Integer3SUM**, $w = \Omega(\log n)$ is the machine word size, and $U \leq 2^w$ is the size of the universe.

(aka APSP-Hardness); see for examples [2, 3, 42]. The second is that CNF-SAT takes $\Omega(2^{(1-o(1))n})$ time. The latter is often referred to as the *Strong Exponential Time Hypothesis* (SETH) [32, 33]. Assuming this hypothesis, Bringmann [12] showed that the Fréchet distance problem cannot be solved in $O(n^{2-\Omega(1)})$ time, and Backurs and Indyk [7] showed a similar result for the famous Edit Distance problem. A natural question is whether any of these conjectures (3SUM, SETH, APSP) are in fact equivalent, or whether they all derive from a basic unifying hypothesis. At the current state of knowledge, there is no strong relationship between any pair of these problems, so it may be possible that the any one of them could be true or false, independently of the status of the others. A recent breakthrough by Carmosino, Gao, and Impagliazzo [15] provides evidence that such a relationship is *unlikely*, based on a variant of SETH, named NSETH; see [15] for details.

1.1 Our Results

Our results are summarized in the following theorems. Refer to Figure 1.1 for a summary of previous work and our results.

Theorem 1.1. *3SUM can be solved deterministically in $O(n^2 \log \log n / \log n)$ time.*

Theorem 1.2. *The randomized 4-linear decision tree complexity of 3SUM is $O(n^{3/2})$.*

Theorem 1.3. *The randomized $(2k - 2)$ -linear decision tree complexity of k -LDT is $O(n^{k/2})$, for any odd $k \geq 3$.*

Theorem 1.1 improves the $O(n^2(\log \log n / \log n)^{2/3})$ -time deterministic algorithm for 3SUM, given in [29], and of their randomized bound.⁴ This also gives the fastest deterministic algorithm for Integer3SUM, and for Convolution3SUM⁵ (see [29] for previous bounds for the latter). Theorem 1.2 and Theorem 1.3 should be compared with the respective $O(n^{3/2} \sqrt{\log n})$ -depth and $O(n^{k/2} \sqrt{\log n})$ -depth decision trees in [29].

Remark 1.4. After completion of our work on (an initial version of) this paper, it was brought to our attention that an independent result of Freund [27] is now available online, and achieves the result of Theorem 1.1, by using similar ideas as those presented in Section 4. These ideas extend (and are based on) the techniques used by Grønlund and Pettie [29]. Nevertheless, Theorems 1.2 and 1.3 do not appear in [27], and the techniques used to derive them are different (and in our opinion, somewhat more sophisticated and novel) than those used for Theorem 1.1.

1.2 An Overview

Our algorithms extend the techniques of Grønlund and Pettie [29], which borrow liberally from the techniques used in Fredman’s 1976 work on the complexity of sorting $X + Y$ [25], and on the decision tree complexity of $(\min, +)$ -matrix multiplication [26], a fundamental problem in P, known for being computationally equivalent to the APSP (all pairs shortest paths) problem in directed graphs with arbitrary real edge weights.

Grønlund and Pettie [29] combined Fredman’s ideas with a geometric domination technique introduced by Chan [16], who used it to develop a faster algorithm for $(\min, +)$ -matrix multiplication. Bremner *et al.* [11] presented some novel ideas on how this geometric domination technique can be usefully implemented to improve other algorithms. We extend these ideas and techniques to obtain our improved algorithm. We also introduce, for the proofs of Theorems 1.1 and 1.3, a special variant of fractional cascading (see [18, 19] and below).

2 Methods and Lemmas

We give an overview of the techniques discussed above. This includes works by Fredman [25] and Chan [16]. In some of our results we will use a special randomized variant of *fractional cascading* (Chazelle and Guibas [18, 19]), which we also review here. This will later allow us to review the results of Grønlund and Pettie [29] and to obtain our new results.

⁴As in [29], the model of computation is a simplified Real RAM model. Real numbers are subject to only two unit-time operations: addition and comparison. In all other respects the machine behaves like a $w = O(\log n)$ -bit word RAM with the standard repertoire of unit-time AC^0 operations: bitwise Boolean operations, left and right shifts, addition, and comparison.

⁵Given a vector $A \in \mathbb{R}^n$, the Convolution3SUM problem is to determine whether there exist indices i, j for which $A(i) + A(j) = A(i + j)$. This problem can easily be reduced to 3SUM; see [29].

Throughout the paper we refer to the trivial (albeit ingenious) observation that $a + b < a' + b'$ iff $a - a' < b' - b$ as *Fredman's trick*. We denote by $[N]$ the first $[N]$ natural numbers succeeding zero $\{1, \dots, [N]\}$, where N may or may not be an integer.

Fredman showed that, given n numbers whose sorted order is one of $\Pi \leq n!$ realizable permutations, they can be sorted using a linear number of comparisons when Π is sufficiently small. Concretely, we have:

Lemma 2.1 (Fredman 1976 [25]). *A list L of n numbers, whose sorted order is one of Π possible permutations, can be sorted with $2n + \log \Pi$ pairwise comparisons.*

Sorting Pairwise Sums and its Geometric Interpretation. Fredman describes the relation between the complexity of hyperplane arrangements and the decision tree complexity of sorting pairwise sums. Grønlund and Pettie [29] use similar arguments in their 3SUM decision tree where they sort pairwise sums. Given lists $A = (a_i)_{i \in [n]}$ and $B = (b_i)_{i \in [n]}$ of distinct real numbers, define the pairwise sum $A + B = \{a_i + b_j \mid i, j \in [n]\}$. The input A, B can be regarded as a point $p = (a_1, \dots, a_n, b_1, \dots, b_n) \in \mathbb{R}^{2n}$. The points in \mathbb{R}^{2n} that agree with a fixed permutation of $A + B$ form a convex cone bounded by the set H of the $\binom{n^2}{2}$ hyperplanes $x_i + y_j - x_k - y_l = 0$, for $i, j, k, l \in [n]$, $(i, j) \neq (k, l)$. The number of possible sorted orders of $A + B$ is therefore bounded by the number of regions (of all dimensions) in the arrangement $\mathcal{A}(H)$ of H . As shown by Buck [13], the number of regions in an arrangement of m hyperplanes in \mathbb{R}^d of dimension $k \leq d$ is at most

$$\binom{m}{d-k} \left(\binom{m-d+k}{0} + \binom{m-d+k}{1} + \dots + \binom{m-d+k}{k} \right).$$

Thus, the number of regions of all dimensions is $O(m^d)$ (where the constant of proportionality is independent of d). Hence, the number of possible sorting permutations of $A + B$ is $O((n^4)^{2n}) = O(n^{8n})$. One can also construct the hyperplane arrangement explicitly in $O(m^d)$ time by a standard incremental algorithm [23]. The following lemma, taken from Grønlund and Pettie [29], extends this analysis by considering only a subset of these hyperplanes, and is therefore an immediate consequence of these observations.

Lemma 2.2. *Let $A = (a_i)_{i \in [n]}$ and $B = (b_i)_{i \in [n]}$ be two lists, each of n real numbers, and let $F \subseteq [n]^2$ be a set of positions in the $n \times n$ grid. The number of realizable orders of $(A + B)|_F := \{a_i + b_j \mid (i, j) \in F\}$ is $O\left(\binom{|F|}{2}^{2n}\right)$, and therefore $(A + B)|_F$ can be sorted with at most $2|F| + 4n \log |F| + O(1)$ comparisons.*

In Lemma 2.2, the case $F = [n]^2$ goes back to Fredman [25], who showed that $O(n^2)$ comparisons suffice to sort $A + B$.

For the algorithms presented and reviewed in this paper, it is important to assume that the elements of the pairwise sum are distinct, and therefore have a unique sorting permutation. When numbers do appear multiple times, a unique sorting permutation can be obtained by breaking ties consistently (see [29] for details).

Bichromatic Dominance Reporting. Given a finite set P of red points and blue points in \mathbb{R}^d , the *bichromatic dominating pairs* problem is to enumerate all the pairs $(p, q) \in P^2$ such that p is red, q is blue, and p dominates q , i.e., p is greater than or equal to q at each of the d coordinates. A natural divide-and-conquer algorithm [36, p. 366] runs in $O(|P| \log^d |P| + K)$ time, where K is the output size. Chan [16] provided an improved bound when d is logarithmic in $|P|$:

Lemma 2.3 (Chan [16]). *Given a finite set $P \subseteq \mathbb{R}^d$ of red and blue points, one can compute all bichromatic dominating pairs $(p, q) \in P^2$ in time $O(c_\epsilon^d |P|^{1+\epsilon} + K)$, where K is the output size, $\epsilon \in (0, 1)$ is arbitrary, and $c_\epsilon = 2^\epsilon / (2^\epsilon - 1)$.*

Throughout the paper, we invoke Lemma 2.3 a large number of times, with $\epsilon = 1/2$, $c_\epsilon \approx 3.42$, and $d = \delta \log n$, where $\delta > 0$ is sufficiently small to make the overall running time of all the invocations dominated by the total output size; see below for details.

Iterative Search and Fractional Cascading. In our decision tree construction for 3SUM, we aim to speed-up binary searches of the same number, in many sorted sets. We will use for this task a special randomized variant of *fractional cascading*, which will be described in Section 5. First, we briefly recall the standard fractional cascading technique, which was introduced by Chazelle and Guibas [18, 19], for solving the *iterative search problem*, defined as follows. Let U be an ordered universe of keys. Define a *catalog* as a finite ordered subset of U . Given a set of k catalogs C_1, C_2, \dots, C_k over U , such that $|C_i| = n_i$ for each $i \in [k]$, and $\sum_{i=1}^k n_i = n$, the iterative search problem is to provide a data structure that supports efficient execution of queries of the form: given a query $x \in U$, return the largest value less than or equal to x in each of the k catalogs.

Fractional cascading lets one preprocess the catalogs in $O(n)$ time, using $O(n)$ storage, and answer iterative search queries in $O(\log n + k)$ time per query. This is essentially optimal in terms of query time, storage size and preprocessing time. The idea is to maintain a sufficient number of pointers across catalogs, so that, once we have the answer c_i to a query in a catalog C_i , we can follow a pointer to an element in C_{i+1} , which is only $O(1)$ indices away from the answer c_{i+1} in C_{i+1} .

In order to obtain optimal query time, the fractional cascading method expands each catalog C_i to an augmented catalog L_i , starting with L_k and proceeding backwards down to L_1 . L_k is the same as C_k , and for each $1 \leq i < k$, L_i is formed by merging C_i with every second element of L_{i+1} . The items in C_i that were not originally in the catalog are marked as synthetic keys. From each synthetic key in C_i we add a bridge (pointer) to the element in L_{i+1} on which it was based. Using these bridges and additional pointers, from each real key to the two consecutive synthetic keys nearest to it, one can follow directly from each element of L_i (real or synthetic) to the elements in L_{i+1} nearest to it, and by construction, the gap between these elements is 2. Thus, given a query number x , after spending $O(\log n)$ time for searching it in L_1 , it takes only $O(1)$ time to locate x in each subsequent catalog, for a total of $O(\log n + k)$ time, as described.

Fractional cascading can also be extended to support a collection of catalogs stored at the vertices of a directed acyclic graph (DAG), and each query searches with some specified element x through the catalogs stored at the nodes of some specified path in the DAG. In more detail, a *catalog graph* is a DAG in which each vertex stores a catalog (ordered list of keys). A query consists of a key x and a path π in the graph, and the goal is to search with x in the catalog of each node of π . When the maximum in/out degree Δ of the catalog graph is constant, fractional cascading can be extended to this scenario, with the same bounds as before (albeit with larger constants of proportionality). Here too each catalog C_v at a node v , is expanded into an augmented catalog L_v , and each L_v passes to its predecessors every 2Δ -th element (instead of every second element in the earlier case, where Δ was 1). See [18, 19] for more details on the construction of the data structure, proof of correctness, and performance analysis.

In our algorithms we will present a special non-standard variant of this method, that lets us preserve the advantages of the other techniques we use.

The Quadratic 3SUM Algorithm. We next give a brief overview of the quadratic-time algorithm. We follow the implementation given by Grønlund and Pettie [29], which is slightly different from the standard approach, but is useful for the explanation of the subquadratic algorithms in [29] and in this paper. For later references, we present the algorithm for the more general three-set version of 3SUM. We recall that in this setup we have three finite sets $A, B, C \subset \mathbb{R}$, and the problem is to determine whether there exist $a \in A, b \in B, c \in C$ such that $a + b + c = 0$. If the sets $A, B, C \subset \mathbb{R}$ are not all of the same size, the problem is named unbalanced 3SUM, as noted earlier.

The algorithm runs over each $c \in C$ and searching for $-c$ in the pairwise sum $A + B$. With a careful implementation, given below, each search takes $O(|A| + |B|)$ time, for a total of $O(|C|(|A| + |B|))$ time. We view $A + B$ as being a matrix whose rows correspond to the elements of A and columns to the elements of B , both listed in increasing order. To help visualizing some steps of the algorithms, we think of the rows arranged in increasing order from top to bottom, and of the columns from left to right, see, e.g., Figure 4.1.

1. Sort A and B in increasing order as $A(0), \dots, A(|A| - 1)$ and $B(0), \dots, B(|B| - 1)$.
2. For each $c \in C$,
 - 2.1. Initialize $\text{lo} \leftarrow 0$ and $\text{hi} \leftarrow |B| - 1$.
 - 2.2. Repeat:
 - 2.2.1. If $-c = A(\text{lo}) + B(\text{hi})$, report witness “ $(A(\text{lo}), B(\text{hi}), c)$ ”.
 - 2.2.2. If $-c > A(\text{lo}) + B(\text{hi})$ then increment lo , otherwise decrement hi .
 - 2.3. Until $\text{lo} = |A|$ or $\text{hi} = -1$.
3. If no witnesses were found report “no witness.”

The correctness easily follows from the fact that each row and column of $A + B$ is sorted in increasing order. Note that when a witness is discovered in Step 2.2.1, the algorithm can stop right there. However, in order to simplify future definitions and explanations, this implementation continues to search for more witnesses. For our purpose, after finding a witness we will always choose to decrement hi . This choice will be made throughout the paper.

Define the *contour* of x , $\text{CONTOUR}(x, A + B)$, ($\text{CONTOUR}(x)$, when the context is clear) to be the sequence of positions (lo, hi) encountered while searching for x in $A + B$ in the preceding algorithm. Lemma 2.4 is straightforward.

Lemma 2.4. *For $x < y \in \mathbb{R}$, $\text{CONTOUR}(x)$ lies fully above $\text{CONTOUR}(y)$; that is, for each $i, i', j \in \{0, \dots, n - 1\}$, if $(i, j) \in \text{CONTOUR}(x)$ and $(i', j) \in \text{CONTOUR}(y)$, then $i \leq i'$.*

By Lemma 2.4 a pair of contours can overlap, but never cross. Moreover, Lemma 2.4 implies a weak total order relation $<$ on the contours, which corresponds to the order between the searched elements, such that $x < y$ iff $\text{CONTOUR}(x) < \text{CONTOUR}(y)$, where the latter relation means that the two contours satisfy the properties stated in the lemma; see Figure 4.1.

3 Grønlund and Pettie’s Subquadratic 3SUM

In this section we give an overview of the recent subquadratic breakthrough of Grønlund and Pettie [29]. In the following sections we show how their techniques can be extended and combined with additional techniques, to yield our improved results.

3.1 Subquadratic Decision Tree

We give an overview of the subquadratic decision tree for 3SUM over a single input set A of size n , taken from [29], resulting in a 4-linear decision tree with depth $O(n^{3/2}\sqrt{\log n})$. This is shown by an algorithm that performs at most $O(n^{3/2}\sqrt{\log n})$ comparisons, where each comparison is a sign test of a linear expression with at most 4 terms.

1. Sort A in increasing order as $A(0), \dots, A(n-1)$. Partition A into $\lceil n/g \rceil$ groups $A_1, \dots, A_{\lceil n/g \rceil}$, each of at most g consecutive elements, where g is a parameter that we will fix later, by setting $A_i := \{A((i-1)g), \dots, A(ig-1)\}$, for each $i = 1, \dots, \lceil n/g \rceil - 1$, where $A_{\lceil n/g \rceil}$ may be smaller. The first and last elements of A_i are $\min(A_i) = A((i-1)g)$ and $\max(A_i) = A(ig-1)$.
2. Sort $D := \bigcup_{i \in [n/g]} (A_i - A_i) = \{a - a' \mid a, a' \in A_i \text{ for some } i\}$.
3. For all $i, j \in [n/g]$, sort $A_{i,j} := A_i + A_j = \{a + b \mid a \in A_i \text{ and } b \in A_j\}$.
4. For k from 1 to n ,
 - 4.1. Initialize $\text{lo} \leftarrow 1$ and $\text{hi} \leftarrow \lceil n/g \rceil$.
 - 4.2. Repeat:
 - 4.2.1. If $-A(k) \in A_{\text{lo,hi}}$, report “solution found” and halt.
 - 4.2.2. If $\max(A_{\text{lo}}) + \min(A_{\text{hi}}) > -A(k)$ then decrement hi , otherwise increment lo .
 - 4.3. Until $\text{lo} = \lceil n/g \rceil + 1$ or $\text{hi} = 0$.
5. Report “no solution” and halt.

This algorithm can be generalized in a straightforward way to solve the (unbalanced) three-set version of 3SUM. For the easy argument concerning the correctness of the algorithm, see [29].

With a proper choice of g , the decision tree complexity of the algorithm is $O(n^{3/2}\sqrt{\log n})$. Step 1 requires $O(n \log n)$ comparisons. By Lemma 2.2, Step 2 requires $O(n \log n + |D|) = O(n \log n + gn)$ comparisons to sort D . By Fredman’s trick, if $a, a' \in A_i$ and $b, b' \in A_j$, $a + b < a' + b'$ holds iff $a - a' < b' - b$, and both sides of this inequality are elements of D . Thus, Step 3 does not require any real input comparisons, given the sorted order on D . For each iteration of the outer loop (in Step 4) there are at most $2\lceil n/g \rceil$ iterations of the inner loop (Step 4.2), since each iteration ends by either incrementing lo or decrementing hi . In Step 4.2.1 we can determine whether $-A(k)$ is in $A_{\text{lo,hi}}$ using binary search, in $\log |A_{\text{lo,hi}}| = O(\log g)$ comparisons. The total number of comparisons is thus $O(n \log n + gn + (n^2 \log g)/g)$, which becomes $O(n^{3/2}\sqrt{\log n})$ when $g = \sqrt{n \log n}$.

3.2 Subquadratic Algorithms

Grønlund and Pettie [29] present two subquadratic algorithms for 3SUM, one is relatively simple, and the second one has slightly faster runtime but is more involved. Both algorithms are based on the decision tree algorithm as just outlined, except that they use a much smaller value of g , in order to make the overall running time subquadratic. We give here a brief overview of the simpler algorithm. Their second algorithm has some common high-level features with our algorithm, presented in the next section, but our algorithm processes the data in a different, simpler, and more efficient manner.

Note that, sorting the set D in the subquadratic decision tree lets one obtain a comparison-efficient way to sort each of $A_{i,j}$. However, the actual running time is even more than quadratic, when all operations are considered. When the boxes $A_{i,j}$ are small enough, Grønlund and Pettie

showed that it is possible to obtain the sorted orders in each of the $(n/g)^2$ boxes, in (all inclusive) subquadratic time.

Specifically, the algorithm enumerates *every* permutation $\pi : [g^2] \rightarrow [g^2]$, where $\pi = (\pi_r, \pi_c)$ is decomposed into row and column functions $\pi_r, \pi_c : [g^2] \rightarrow [g]$, so that $\pi(k) = (\pi_r(k), \pi_c(k))$, for each $k \in [g^2]$. By definition, π is the correct sorting permutation for the box $A_{i,j}$ iff $A_{i,j}(\pi(t)) < A_{i,j}(\pi(t+1))$ for all $t \in [g^2 - 1]$. Since $A_{i,j} = A_i + A_j$ this inequality can also be written $A_i(\pi_r(t)) + A_j(\pi_c(t)) < A_i(\pi_r(t+1)) + A_j(\pi_c(t+1))$. By Fredman's trick this is equivalent to saying that the (red) point p_j dominates the (blue) point q_i , where

$$\begin{aligned} p_j &= (A_j(\pi_c(2)) - A_j(\pi_c(1)), \dots, A_j(\pi_c(g^2)) - A_j(\pi_c(g^2 - 1))) \\ q_i &= (A_i(\pi_r(1)) - A_i(\pi_r(2)), \dots, A_i(\pi_r(g^2 - 1)) - A_i(\pi_r(g^2))) \end{aligned}$$

Invoking $(g^2)!$ times the bichromatic dominance reporting algorithm from Lemma 2.3, we find, for each π , all such dominating pairs, that is, all boxes $A_{i,j}$ sorted by π . Note that, for each pair of indices j, i , there is exactly one invocation of the dominating pairs procedure in which the corresponding points p_j and q_i are such that p_j dominates q_i ; this follows because we assume that all elements of $A_{i,j}$ are distinct (see a previous remark concerning this issue). This is important in order to keep the overall output size subquadratic.

By Lemma 2.3 and the remarks just made, the time to report all red/blue dominating pairs, over all $(g^2)!$ invocations of the procedure, is $O\left((g^2)! c_\epsilon^{g^2-1} (2n/g)^{1+\epsilon} + (n/g)^2\right)$, where the last term is the total size of the outputs (one for each box $A_{i,j}$). For $\epsilon = 1/2$ and $g = \frac{1}{2}\sqrt{\log n / \log \log n}$, the first term turns out to be negligible. The total running time is therefore $O((n/g)^2)$ for dominance reporting, and $O(n^2 \log g/g) = O(n^2 (\log \log n)^{3/2} / (\log n)^{1/2})$ for the binary searches in Steps 4.1–4.3. By Lemma 2.2 and Fredman [25], there are at most $O(g^{8g})$ realizable permutations of $A_{i,j}$ (which is much smaller than $(g^2)!$). Hence, this algorithm can be slightly improved to run in $O(n^2 \log \log n / \sqrt{\log n})$ time, by constructing the arrangement of the hyperplanes (as defined in Section 2) explicitly, extracting from it the relevant permutations, and choosing $g = \Theta(\sqrt{\log n})$.

4 Improved Deterministic Subquadratic 3SUM Algorithm

In the algorithm of Grønlund and Pettie, described above, the boxes $A_{i,j}$ are sorted by using Fredman's trick to transform each permutation into a sequence of $g^2 - 1$ comparisons, which are then resolved by the bichromatic dominance reporting algorithm. Consequently, the space into which these sequences are encoded is of dimension $g^2 - 1$, thus having the $c_\epsilon^{g^2-1}$ factor in the running time of the bichromatic dominance reporting algorithm forced us to use $g = \Theta(\sqrt{\log n})$. In order to use a larger value of g , we want to reduce the dimension of the points. Thus, we want to find a method to sort smaller sets, while still be able to do the binary searches in each box in $O(\log g)$ time.

Fix some $k \in [g^2]$, and let $(l, m) \in [g]^2$ be a point in the $g \times g$ grid, such that $A_{i,j}(l, m)$ is the k -th smallest element in the box $A_{i,j}$. Let $\tau = (\tau_r, \tau_c)$ denote $\text{CONTOUR}(A_{i,j}(l, m))$, and enumerate its elements as $\tau(1), \tau(2), \dots$. Recall that, if $\tau(t+1) = \tau(t) + (0, -1)$ then $A_{i,j}(l, m) \leq A_{i,j}(\tau(t))$, otherwise; if $\tau(t+1) = \tau(t) + (1, 0)$ then $A_{i,j}(l, m) > A_{i,j}(\tau(t))$. The contour starting position is $(\tau_r(0), \tau_c(0)) = (1, g)$, and it ends at the first t^* for which $\tau(t^*) = (g+1, \cdot)$ or $\tau(t^*) = (\cdot, 0)$. Recall that a pair of contours $\text{CONTOUR}(x)$ and $\text{CONTOUR}(y)$ in $A_{i,j}$ may overlap, but can never cross; see Lemma 2.4.

Let $\tau'(A_{i,j}(l, m)) = (\tau'(0), \tau'(1), \dots, \tau'(t_{\tau'})) \subseteq \tau = \text{CONTOUR}(A_{i,j}(l, m))$ be the *partial contour* of τ , defined as the subsequence of positions of τ at which we chose to go down (i.e., increment lo

372	389	407	439	454	480	534	609	635	655
397	414	432	464	479	505	559	634	660	680
420	437	455	487	502	528	582	657	683	703
442	459	477	509	524	550	604	679	705	725
478	495	513	545	560	586	640	715	741	761
500	517	535	567	582	608	662	737	763	783
523	540	558	590	605	631	685	760	786	806
548	565	583	615	630	656	710	785	811	831
594	611	629	661	676	702	756	831	857	877
627	644	662	694	709	735	789	864	890	910

Figure 4.1: The sky-blue colored entries form $\text{CONTOUR}(710)$, and the purple colored ones form $\text{CONTOUR}(558)$; A shared cell is shown in green. The lighter colors (light purple and light sky-blue) depict their *partial contour*, that is, the positions of the contours where we chose to go down. All the elements in the matrix whose values are in $[558, 710)$ are enclosed between these two contours, excluding the partial contour of 558 and including the partial contour of 710.

in the quadratic algorithm). The sequence $\tau'(A_{i,j}(l, m))$ is of length at most g , since it contains at most one element of each row (at which we go down, by incrementing l); see Figure 4.1.

Since the rows of $A_{i,j}$ are sorted, each position $(a, b) \in \tau'(A_{i,j}(l, m))$ satisfies

$$A_{i,j}(a, b') < A_{i,j}(l, m) \quad \text{for every } b' \leq b. \quad (1)$$

$$A_{i,j}(a, b'') \geq A_{i,j}(l, m) \quad \text{for every } b'' > b. \quad (2)$$

Thus, $\tau'(A_{i,j}(l, m))$ partitions $A_{i,j}$ into two sets: $A_{(i,j)L} \subset (-\infty, A_{i,j}(l, m))$ consists of the elements at positions in $\{(a, b') \mid (a, b) \in \tau'(A_{i,j}(l, m)) \text{ and } b' \leq b\}$, and $A_{(i,j)R} \subset [A_{i,j}(l, m), \infty)$ consists of the elements at positions in $\{(a, b'') \mid (a, b) \in \tau'(A_{i,j}(l, m)) \text{ and } b'' > b\}$. Rows succeeding the last row of $\tau'(A_{i,j}(l, m))$ are fully contained in $A_{(i,j)R}$. By construction, $A_{(i,j)L}$ is the set of all elements in $A_{i,j}$ that are smaller than the k -th smallest element $A_{i,j}(l, m)$, so the considerations just made, provide the structure of this set. See Figure 4.1 for an illustration.

Each partial contour τ' is thus a sequence of positions in $A_{i,j}$ such that (i) the rows containing these positions form a contiguous subsequence, starting from the first row of $A_{i,j}$, (ii) each row in this subsequence has exactly one entry of τ' , and (iii) the sequence of columns of the entries of τ' is weakly monotone decreasing: if (a, b) and $(a + 1, b')$ are in τ' then $b' \leq b$. Any sequence τ' that satisfies properties (i)–(iii) is called a *valid partial contour*. Note that a valid partial contour depends only on the positions of the contour in the box $A_{i,j}$, and not on the actual values of the entries of $A_{i,j}$.

Let μ' be some valid partial contour, as just defined, over a $[g] \times [g]$ position set, such that the sum of the column indices of positions in μ' is exactly k . Write $\mu' = (\mu'_r, \mu'_c)$, as was done for permutations above, so that μ'_r gives the row indices of the elements of μ' , and μ'_c gives their column indices. Denote by $t' \leq g$ the number of positions in μ' .

For given indices ℓ, m , we can determine, using (1) and (2), whether $\mu' = \tau'(A_{i,j}(l, m))$, by testing, for each $t \in [t']$, whether $A_{i,j}(\mu'(t)) < A_{i,j}(l, m)$ and $A_{i,j}(\mu'(t) + (0, 1)) > A_{i,j}(l, m)$, except for the t_0 for which $A_{i,j}(\mu(t_0)) = A_{i,j}(l, m)$, since then the second inequality becomes an equality; this takes at most $2t' - 2$ comparisons. By Fredman's trick, and since $A_{i,j} = A_i + A_j$, this can be

restated, that $\mu' = \tau'(A_{i,j}(l, m))$ iff the (red) point p_j dominates the (blue) point q_i , where

$$\begin{aligned} p_j &= (\dots, A_j(m) - A_j(\mu'_c(t)), A_j(\mu'_c(t + (0, 1)) - A_j(m)), \dots) \\ q_i &= (\dots, A_i(\mu'_r(t)) - A_i(l), A_i(l) - A_i(\mu'_r(t)), \dots), \end{aligned} \quad (3)$$

where the $2t' - 2$ coordinates are indexed in pairs by $t \in [t'] - \{t_0\}$.

We regard each box $A_{i,j}$ as being partitioned into $h = g \log g$ sets $A_{(i,j)1}, \dots, A_{(i,j)h}$, each of size at most $s = g/\log g$, such that for $k \in [h]$, $A_{(i,j)k}$ is the set of all elements that are at least the $(k-1)s$ -smallest element, and smaller than the ks -smallest element in $A_{i,j}$. Our goal is to compute, for each box $A_{i,j}$, the positions of the elements of the sets $A_{(i,j)1}, \dots, A_{(i,j)h}$, and the correct sorting permutation of each of them, as well as to determine, for each $k \in [h]$, the position of the ks -smallest element in $A_{i,j}$.

Fix $k \in [h]$. We enumerate all the pairs of realizable valid partial contours $\mu'_{(k-1)s}, \mu'_{ks}$, such that (i) $\mu'_{(k-1)s}$ lies to the left and above μ'_{ks} , and (ii) the sums of the column indices of their entries are $(k-1)s$ and ks , respectively. Let S_k be the set of positions enclosed between the two partial contours $\mu'_{(k-1)s}$ and μ'_{ks} , excluding $\mu'_{(k-1)s}$ and including μ'_{ks} . For each $A_{i,j}$, we want to identify the pair $(\mu'_{(k-1)s}, \mu'_{ks})$, for which $\mu'_{(k-1)s}$ and μ'_{ks} are the partial contours of the $(k-1)s$ -smallest and the ks -smallest elements of $A_{i,j}$, respectively. Thus S_k is the set of the s positions of the elements of $A_{i,j}$ that are larger or equal to the $(k-1)s$ -smallest element and smaller than the ks -smallest element. These are the positions of the set $A_{(i,j)k}$. See Figure 4.1 for an illustration.

It is easily seen that there are at most 2^{4g} pairs of sequences $(\mu'_{(k-1)s}, \mu'_{ks})$, and there is only one unique pair of valid partial contours $(\mu'_{(k-1)s}, \mu'_{ks})$ that satisfy all the above requirements for a specific box $A_{i,j}$, as there is only one $(k-1)s$ -smallest element and only one ks -smallest element in $A_{i,j}$ (assuming, as above, that all the elements of $A_{i,j}$ are distinct). We enumerate all pairs of positions $P_1, P_2 \in [g]^2$, such that $P_1 \in \mu_{(k-1)s}$ and $P_2 \in \mu_{ks}$ (recall that μ' is a partial contour of some contour μ , where μ is uniquely determined from μ' , see Figure 4.1). There are at most $(2g)^2 = 4g^2$ such positions. We also enumerate every realizable permutation $\pi : [s] \rightarrow S_k$ of the elements at positions in S_k (where, for each $A_{i,j}$, we want to identify the permutation that sorts its elements at the positions of S_k). The number of permutations is bounded trivially by $s! = (g/\log g)!$.

We now extend the points defined in (3), to make them encode additional information, as follows. For every tuple $(P_1, P_2, \mu'_{(k-1)s}, \mu'_{ks}, \pi)$, we create red points $\{p_j\}_{j \in [n/g]}$ and blue points $\{q_i\}_{i \in [n/g]}$ in $\mathbb{R}^{4(t'-1)+s-1}$, such that p_j dominates q_i iff (i) $\mu'_{(k-1)s} = \tau'(A_{i,j}(P_1))$, (ii) $\mu'_{ks} = \tau'(A_{i,j}(P_2))$, and (iii) π is the unique sorting permutation of the portion of $A_{i,j}$ with indices in S_k . The first $4t' - 4$ coordinates encode the correctness of $\mu'_{(k-1)s}$ and μ'_{ks} (as in (3), using the positions P_1, P_2 as those defining the respective contours), and the last $s - 1$ coordinates encode the correctness of π , as in Section 3.2 but for a permutation of size at most $s = g/\log g$. We do this $h = g \log g$ times, for each $k \in [h]$.

According to Lemma 2.3, the overall time to report all bichromatic dominating pairs is

$$O\left(h \cdot 2^{4g} g^2 s! c_\epsilon^{4(g-1)+s-1} (n/g)^{1+\epsilon} + h(n/g)^2\right).$$

The second term is the output size, because for each of the $(n/g)^2$ boxes $A_{i,j}$, there will be exactly h dominating pairs, one for each pair of consecutive partial contours, as above. By fixing $\epsilon = 1/2$ and $g = d \log n$ with a small enough d , the first term will be negligible and the runtime will be dominated by the output size $O(h(n/g)^2) = O(n \log g/g) = O(n^2 \log \log n / \log n)$.

We can now search an element x in a box $A_{i,j}$, in $O(\log g)$ time. We first do a binary search, in $O(\log g)$ time, over the h positions storing the ks -smallest element of $A_{i,j}$, for $k \in [h]$ (we

have already computed their positions, and, by definition, they are already sorted). This will give us a single set $A_{(i,j)k}$ that can possibly contain x . Then we do another binary search in $A_{(i,j)k}$, also in $O(\log g)$ time, as we already found its sorting permutation earlier. (Note that each such permutation π is of length at most $g/\log g$, and of values from $[g]^2$. Thus, by our earlier choice of g , π can be stored in a machine word of size $O(\log n)$, and be accessed in $O(1)$ time.) Each element $-A(k)$ is being searched in at most $2\lceil n/g \rceil$ boxes (as in Steps 4.1–4.3 of Grønlund and Pettie’s decision tree, described in Section 3). Hence, the total running time of the algorithm is $O(n^2 \log g/g) = O(n^2 \log \log n / \log n)$ *deterministic* time. This proves Theorem 1.1.

5 Randomized Decision Trees for 3SUM, k -SUM, and k -LDT

In this section we show that the randomized decision tree complexity of 3SUM is $O(n^{3/2})$, and more generally, that the randomized decision tree complexity of k -LDT is $O(n^{k/2})$, for any odd $k \geq 3$. This bound removes the $O(\sqrt{\log n})$ factor in Grønlund and Pettie’s decision tree bound. We show this results by giving a randomized algorithm that constructs a $(2k-2)$ -linear decision tree whose expected depth is $O(n^{k/2})$.

To make the presentation more concise, we present it for the variant where we have three different sets A, B, C of n real numbers each, and we want to determine whether there exist $a \in A$, $b \in B$, $c \in C$, such that $a + b + c = 0$.

As in previous sections, we partition each of the sorted sets A and B into $\lceil n/g \rceil$ blocks, each consisting of g consecutive elements, denoted by $A_1, \dots, A_{\lceil n/g \rceil}$, and $B_1, \dots, B_{\lceil n/g \rceil}$, respectively. As above, but with a slightly different notation, we consider the $n \times n$ matrix $M = M^{AB}$, whose rows (resp., columns) are indexed by the (sorted) elements of A (resp., of B), so that $M(k, \ell) = a_k + b_\ell$, for $k, \ell \in [n]$. The partitions of A and of B induce, as before, a partition of M into n^2/g^2 boxes $M_{i,j}$, for $i, j \in [\lceil n/g \rceil]$, where $M_{i,j}$ is the portion of M with rows in A_i and columns in B_j .

Fredman’s trick allows us to sort all the boxes $M_{i,j}$ with $O(ng)$ comparisons. Since the problem is fully symmetric in A, B, C , we can also define analogous matrices M^{AC} and M^{BC} , constructed in the same manner for the pairs A, C and B, C , respectively, partition each of them into n^2/g^2 boxes, and obtain the sorted orders of all the corresponding boxes, with $O(ng)$ comparisons.

The crucial (costliest) step in Grønlund and Pettie’s algorithm, which we are going to improve, is the searches of the elements of $-C$ in M^{AB} . For each $c \in C$, let $\sigma(c) = \text{CONTOUR}(-c)$ denote the staircase path contour of $-c$, as defined before Lemma 2.4. The length of $\sigma(c)$ is thus at most $2n$. Each of the paths $\sigma(c)$ visits some (at most $2\lceil n/g \rceil$) of the boxes $M_{i,j}$, and the index pairs (i, j) of these boxes also form a staircase pattern, as in the preceding sections. The number of boxes that a contour $\sigma(c)$ visits is at most $2\lceil n/g \rceil$. For each $c \in C$, the sequence of boxes that $\sigma(c)$ visits can be obtained by invoking (an appropriate variant of) Step 4 of the algorithm in Section 3.1, excluding the binary search in Step 4.2.1. The total running time of this step, over all $c \in C$, is $O(n^2/g)$.

The paths $\sigma(c)$, being contours, have the structure given in Lemma 2.4, including the weak total order $<$ between them. As a corollary, we obtain:

Claim: For each box $M_{i,j}$, let $C_{i,j}$ denote the set of elements of C whose paths $\sigma(c)$ traverse $M_{i,j}$. Then $C_{i,j}$ is a contiguous subsequence of C .

Put $\kappa_{i,j} := |C_{i,j}|$. Then we clearly have $\sum_{i,j \in [\lceil n/g \rceil]} \kappa_{i,j} = O(n^2/g)$. That is, the average number of elements of c that visit a box is $O(g)$, and, for each box, these elements form a contiguous subsequence of C , as just asserted. Let $C_{i,j}^*$ denote the contiguous sequence of indices in C of the elements of $C_{i,j}$. That is, $C_{i,j} = \{c_\ell \mid \ell \in C_{i,j}^*\}$. With all these observations, we next proceed to derive the mechanism by which, for each box $M_{i,j}$, we can efficiently search in $M_{i,j}$ with the (negations of the) $\kappa_{i,j}$ corresponding elements of $C_{i,j}$.

We apply a special variant of fractional cascading. The twist is in the way in which we construct the augmented catalogs. Note that in each box $M_{i,j}$, we have g^2 elements of the form $a_k + b_\ell$, but only $2g$ indices k, ℓ . We want to sample elements from a box, and then copy and merge them into its neighbor boxes. However, in order to be able to use Fredman's trick, we have to preserve the property that the number of element-indices (rows and columns) in each box stays $O(g)$ (unlike a naive implementation of fractional cascading, where it is enough that each box be of size $O(g^2)$).

Thus, we sample elements from A (row elements) and elements from B (column elements) separately. We construct augmented sets $A'_1, \dots, A'_{\lceil n/g \rceil}$. Starting with $A'_{\lceil n/g \rceil} = A_{\lceil n/g \rceil}$, we sample each element in $A'_{\lceil n/g \rceil}$ with probability $p = \frac{1}{4}$. Each sampled element is copied and merged with $A_{\lceil n/g \rceil - 1}$, and we denote by $A'_{\lceil n/g \rceil - 1}$ the new augmented set. Then we sample each element from $A'_{\lceil n/g \rceil - 1}$ with the same probability p , copy and merge the sampled elements with $A_{\lceil n/g \rceil - 2}$, obtaining $A'_{\lceil n/g \rceil - 2}$, and continue this process until the augmented set A'_1 is constructed. Similarly, we construct the augmented sets $B'_1, \dots, B'_{\lceil n/g \rceil}$, but we do it in the opposite direction, starting from $B'_1 = B_1$ and ending with $B'_{\lceil n/g \rceil}$. Clearly, similar to standard fractional cascading, the expected size of each of the augmented sets is $O(g)$, as the expected numbers of additional elements placed in each box form a convergent geometric series. Now we sort

$$D_{A'} = \bigcup_{i \in [n/g]} (A'_i - A'_i) = \{a - a' \mid a, a' \in A'_i \text{ for some } i\}.$$

In each $A'_i - A'_i$, the expected number of elements $a_k - a_{k'}$ is $O(g^2)$, and the expected number of element indices k, k' is only $O(g)$. Thus, by Lemma 2.2, we can sort $D_{A'}$ with expected $O(ng)$ comparisons. Similarly, we sort $D_{B'} = \bigcup_{j \in [n/g]} (B'_j - B'_j)$ with the same expected number of comparisons. Then, we form the union $D' = D_{A'} \cup D_{B'}$ and obtain its sorted order by merging $D_{A'}$ and $D_{B'}$. This costs additional expected $O(ng)$ comparisons. By Fredman's trick, from the sorted order of D' , we can obtain the sorted order of the augmented boxes $A'_i + B'_j$, for each $i, j \in [n/g]$, without further comparisons.

With these augmentations of the row and column blocks, the matrix M^{AB} itself is now augmented, such that each modified box $M_{i,j} = A'_i + B'_j$ receives some fraction of the rows from the box $M_{i+1,j}$ below it, and a fraction of the columns from the box $M_{i,j-1}$ to its left. Each box $M_{i,j}$ corresponds to a vertex in the catalog graph, and it has (at most) two outgoing edges, one to the vertex that corresponds to $M_{i+1,j}$ and one to the vertex that corresponds to $M_{i,j-1}$ (it also has at most two incoming edges). Clearly this is a DAG with maximum in/out degree $\Delta = 2$, which is why we sampled $\frac{1}{2\Delta} = \frac{1}{4}$ of the rows/columns in each step. We complete the construction of this special fractional cascading data structure, by adding the appropriate pointers, similar to what is done in a standard implementation of fractional cascading (see Section 2). This does not require any further comparisons, since the pointers from synthetic keys (the sampled elements) to real keys, and pointers from real keys to synthetic keys, depend only on the sorted order of the augmented sets $M_{i,j}$, which we already computed. So the overall expected number of comparisons needed to construct this data structure is still $O(ng)$.

Consider now the search with $-c$, for some $c \in C$. Assume that the search has just visited some box $M_{i',j'}$, and now proceeds to search in box $M_{i,j}$. Thus, either $(i, j) = (i' + 1, j')$ or $(i, j) = (i', j' - 1)$. Assume, without loss of generality, that $(i, j) = (i' + 1, j')$; a symmetric argument applies when $(i, j) = (i', j' - 1)$, using columns instead of rows. In this case, the fractional cascading mechanism has sampled, in a random manner, an expected quarter of the rows of (the already augmented) $M_{i,j}$ and has sent them to $M_{i',j'} = M_{i-1,j}$. The output of the search at $M_{i-1,j}$, if $-c$ was not found there, includes two pointers to the largest element ξ^- of $M_{i,j}$ that is smaller

60	70	80	90	100	110	120	130	140	150
160	170	180	190	200	210	220	230	240	250
260	270	280	290	300	310	320	330	340	350

Figure 5.1: An expensive step in the fractional cascading search: Assume that only the first and third rows (appearing in gray) are sent to the preceding box (above the current one), and that we search with $-c = 205$. The previous search locates $-c$ between $\xi^- = 150$ and $\xi^+ = 260$, say, and now we have to examine the entire second row to locate $-c$ in the current box.

than $-c$, and to the smallest element ξ^+ of $M_{i,j}$ that is larger than or equal to $-c$. We need to go over the elements in the sorted order of $M_{i,j}$ that lie between ξ^- and ξ^+ , and locate $-c$ among them. If we do not find it, we get the two consecutive elements that enclose $-c$, retrieve from them two corresponding pointers to a pair of elements in the next box to be searched, that enclose $-c$ between them, and continue the fractional cascading search in the next box, in between these elements.

The main difficulty in this approach is that the number of elements of $M_{i,j}$ between ξ^- and ξ^+ might be large, because there might be many elements between ξ^- and ξ^+ in rows that we did not sample, and then we have to inspect them all, slowing down the search.

Concretely, in this case we sample, in expectation, a quarter of the rows of $M_{i,j}$ (recall that, we actually sample the rows from an augmented box that has already received data from previous boxes, but let us ignore this issue for now). Collectively, these rows contain (in expectation) $\Theta(g^2)$ elements of $M_{i,j}$, but we have no good control over the size of the gaps of non-sampled elements between consecutive pairs of sampled ones. This is because there might be rows that we did not sample which contain many elements between ξ^- and ξ^+ , and searching through such large gaps could slow down the procedure considerably. See Figure 5.1 for an illustration. (For a normal fractional cascading, this would not be an issue, but here the peculiar and implicit way in which we sample elements has the potential for creating this problem.)

We handle this problem as follows. Consider any gap of non-sampled elements of $M_{i,j}$ between a consecutive pair $\xi^- < \xi^+$ of sampled ones. We claim that the expected number of rows to which these elements belong is 12. Indeed, the probability to have k distinct rows in such a gap is $(3/4)^k$, which follows since each row is sampled *independently* with probability $1/4$. Hence, the expected row-size of a gap is $\sum_{k \geq 0} k(3/4)^k = 12$, as claimed. Denote this expected value as $\beta = 12$. In other words, for each $c \in C_{i,j}$, let R_c be the set of rows that show up in the gap between the corresponding elements ξ^- and ξ^+ for c . The overall expected size $\sum_{c \in C_{i,j}} |R_c|$ is thus $\beta|C_{i,j}|$.

Fix a box $M_{i,j}$. For each $\ell \in C_{i,j}^*$ and for each $k \in R_{c_\ell}$, we need to locate $-c_\ell$ among the elements in row k of $M_{i,j}$. That is, we need to locate $-c_\ell$ among the elements of the set $a_k + B'_j$. This however is equivalent to locating $-a_k - c_\ell$ among the elements of B'_j .

We therefore collect the set S of all the sums $-a_k - c_\ell$, for $\ell \in C_{i,j}^*$ and $k \in R_{c_\ell}$, and recall that in expectation we have $|S| = O(|C_{i,j}|)$. The crucial observation is that we already (almost) know the order of these sums. To make this statement more precise, partition, in the usual manner, the sorted sequence C into $\lceil n/g \rceil$ blocks $C_1, C_2, \dots, C_{\lceil n/g \rceil}$, each consisting of g consecutive elements in the sorted order. As mentioned earlier, a symmetric application of Fredman's trick allows us to obtain the sorted order of each box of the form $A'_i + C_j$, using a total of O/ng comparisons.

The number of (consecutive) blocks C_s of C that overlap $C_{i,j}$ is $t_{i,j} \leq \lceil \kappa_{i,j}/g \rceil + 2$. Moreover, each sum in S belongs to $-(A'_i + C_s)$ for one of these $t_{i,j}$ blocks. Since each of these sets is already sorted, we extract from them (with no extra comparisons) the elements of S as the union of $t_{i,j}$ sorted sequences $S_{i,s}$, where $S_{i,s} \subset -(A'_i + C_s)$ for each s . Arguing as above, the expected size of

$S_{i,s}$ is $\beta|C_s| = O(g)$. We now merge each of the sorted sequences $S_{i,s}$ with B'_j , using an expected $O(g)$ comparisons for each merge. As a result, each sum $-a_i - c_\ell$ is located between two consecutive elements $b_{i,\ell}^- < b_{i,\ell}^+$ of B'_j . In other words, for each $c_\ell \in C_{i,j}$, we have at most $|R_{c_\ell}|$ candidates for being the largest element of $M_{i,j}$ that is smaller than $-c_\ell$ (these are the elements $a_i + b_{i,\ell}^-$, for $i \in R_{c_\ell}$), and we select the largest of them, requiring no comparisons, as these are all elements of the already sorted $A'_i + B'_j$. In the same manner, we find the smallest element of $M_{i,j}$ that is larger than $-c_\ell$. Having found these two elements, we can proceed to search $-c_\ell$ in the next box, using the appropriate pointers created by the fractional cascading mechanism (see Section 2).

The overall number of merges is

$$\sum_{i,j \in [n/g]} t_{i,j} \leq \sum_{i,j \in [n/g]} (\kappa_{i,j}/g + 2) = O(n^2/g^2),$$

and each of them costs $O(g)$ expected comparisons, for a total of $O(n^2/g)$ expected comparisons. Thus, the overall number of expected comparisons is $O(n^2/g + n \log g + n/g)$, which is $O(n^{3/2})$, when $g = \sqrt{n}$. This proves Theorem 1.2.

Remark 5.1. At the moment, we do not see how to extend our algorithm from Section 4 to turn this randomized decision tree into a more efficient randomized algorithm for 3SUM. Even when trying to plug it into the simpler (slower) algorithm from Section 3.2 (where we enumerate all realizable permutations of positions in $[g]^2$), due to the fractional cascading phase, each box may now be of different size than the others. Thus, for some permutation π , if p_j dominates q_i and p_j also dominates q_{i+1} , it cannot be that π is the sorting permutation of both $M_{i,j}$ and $M_{i+1,j}$, as these boxes are of different sizes, and thus, they have different sorting permutations. Hence, the total output size from the invocations of the dominance reporting algorithm could be too large.

Linear Degeneracy Testing. The standard algorithm for k -variate linear degeneracy testing (k -LDT) for odd $k \geq 3$, is based on a straightforward reduction to an instance of unbalanced 3SUM, where $|A| = |B| = n^{(k-1)/2}$ and $|C| = n$; see [5] and [29]. The analysis of this section also applies for unbalanced 3SUM, and directly implies that it can be solved by using an expected number of

$$O(g(|A| + |B| + |C|) + |C|((|A| + |B|)/g + \log g))$$

comparisons, where the first term is the cost of sorting the blocks of (the augmented) M^{AB} , M^{AC} , and M^{BC} , and where the second term is the cost of the fractional cascading searches. We have $|A| = |B| = n^{(k-1)/2}$, $|C| = n$, so by choosing $g = \sqrt{n}$, this becomes $O(n^{k/2})$. Thus, the randomized decision tree complexity of k -LDT (and thus of k -SUM) is $O(n^{k/2})$, for any odd $k \geq 3$, as stated in Theorem 1.3.

6 Discussion

Our results strengthen the doubts (originally raised due to the results in [29]) on the optimality of many algorithms for 3SUM-Hard problems. It seems that in terms of logarithmic factors, our results approach the bottleneck of the mechanisms that we use. Thus any further improvement, of at least some power of a logarithmic factor, is likely to require a novel algorithmic machinery.

In light of the results presented in this paper, it is natural to ask whether similar improvements could be made for some well studied 3SUM-Hard problems. For example, can we obtain a sub-quadratic algorithm for determining whether there are three collinear points in a set of n points in the plane? Can we determine whether all elements in $X + Y$ are distinct (or can we sort $X + Y$) in

$o(n^2 \log n)$ time? In our 3SUM algorithm we showed that we can obtain the sorting permutations of boxes of size $\log n \times \log n$, that comprise $X + Y$, in $O(n^2 \log \log n / \log n)$ time. (For each box $A_{i,j}$, we obtain the sorting permutations of the disjoint subsets $A_{(i,j)1}, \dots, A_{(i,j)h}$ that comprise it. The concatenation of these h permutations, in this order, gives the sorting permutation of the box $A_{i,j}$.) Is it possible to pay an additional $o((\log n)^2 / \log \log n)$ factor and determine whether all elements in $X + Y$ are distinct? These are prominent long-standing open problems, and the results in this paper, on top of those in [29], might further motivate the efforts of tackling them.

References

- [1] A. Abboud and K. Lewi. Exact weight subgraphs and the k -sum conjecture. In *Proceedings of the 40th Int'l Colloquium on Automata, Languages and Programming (ICALP)*, pages 1–12, 2013.
- [2] A. Abboud, V. Vassilevska Williams, and H. Yu. Matching triangles and basing hardness on an extremely popular conjecture. In *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC)*, pages 41–50, 2015.
- [3] A. Abboud and V. V. Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *Proceedings of the 55th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 434–443, 2014.
- [4] O. Aichholzer, F. Aurenhammer, E. D. Demaine, F. Hurtado, P. Ramos, and J. Urrutia. On k -convex polygons. *Comput. Geom.*, 45(3):73–87, 2012.
- [5] N. Ailon and B. Chazelle. Lower bounds for linear degeneracy testing. *J. ACM*, 52(2):157–171, 2005.
- [6] A. Amir, T. M. Chan, M. Lewenstein, and N. Lewenstein. On hardness of jumbled indexing. In *Proceedings of the 41st Int'l Colloquium on Automata, Languages, and Programming (ICALP)*, pages 114–125, 2014.
- [7] A. Backurs and P. Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). In *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC)*, pages 51–58, 2015.
- [8] I. Baran, E. D. Demaine, and M. Pătraşcu. Subquadratic algorithms for 3SUM. *Algorithmica*, 50(4):584–596, 2008.
- [9] G. Barequet and S. Har-Peled. Polygon containment and translational min-Hausdorff-distance between segment sets are 3SUM-hard. *Int. J. Comput. Geometry Appl.*, 11(4):465–474, 2001.
- [10] A. Björklund, R. Pagh, V. V. Williams, and U. Zwick. Listing triangles. In *Proceedings of the 41st Int'l Colloquium on Automata, Languages, and Programming (ICALP)*, pages 223–234, 2014.
- [11] D. Bremner, T. M. Chan, E. D. Demaine, J. Erickson, F. Hurtado, J. Iacono, S. Langerman, M. Pătraşcu, and P. Taslakian. Necklaces, convolutions, and $X + Y$. *Algorithmica*, 69:294–314, 2014.

- [12] K. Bringmann. Why walking the dog takes time: Fréchet distance has no strongly subquadratic algorithms unless SETH fails. In *Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 661–670, 2014.
- [13] R. C. Buck. Partition of space. *Amer. Math. Monthly*, 50:541–544, 1943.
- [14] A. Butman, P. Clifford, R. Clifford, M. Jalsenius, N. Lewenstein, B. Porat, E. Porat, and B. Sach. Pattern matching under polynomial transformation. *SIAM J. Comput.*, 42(2):611–633, 2013.
- [15] M. Carmosino, J. Gao, R. Impagliazzo, I. Mikhailin, R. Paturi, and S. Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. *Electronic Colloquium on Computational Complexity (ECCC)*, 22:148, 2015.
- [16] T. M. Chan. All-pairs shortest paths with real weights in $O(n^3/\log n)$ time. *Algorithmica*, 50(2):236–243, 2008.
- [17] T. M. Chan and M. Lewenstein. Clustered integer 3SUM via additive combinatorics. In *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC)*, pages 31–40, 2015.
- [18] B. Chazelle and L. J. Guibas. Fractional cascading: I. A data structuring technique. *Algorithmica*, 1(2):133–162, 1986.
- [19] B. Chazelle and L. J. Guibas. Fractional cascading: II. Applications. *Algorithmica*, 1(2):163–191, 1986.
- [20] S. Chechik, D. Larkin, L. Roditty, G. Schoenebeck, R. E. Tarjan, and V. V. Williams. Better approximation algorithms for the graph diameter. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1041–1052, 2014.
- [21] K.-Y. Chen, P.-H. Hsu, and K.-M. Chao. Approximate matching for run-length encoded strings is 3SUM-Hard. In G. Kucherov and E. Ukkonen, editors, *Combinatorial Pattern Matching*, volume 5577 of *Lecture Notes in Computer Science*, pages 168–179. Springer-Verlag Berlin-Heidelberg, 2009.
- [22] E. D. Demaine, J. S. B. Mitchell, and J. O’Rourke. The open problems project. <http://cs.smith.edu/~orourke/TOPP/P41.html>. Accessed: 2015-10-28.
- [23] H. Edelsbrunner, J. O’Rourke, and R. Seidel. Constructing arrangements of lines and hyperplanes with applications. *SIAM J. Comput.*, 15(2):341–363, 1986.
- [24] J. Erickson. Bounds for linear satisfiability problems. *Theor. Comput. Sci*, 8:388–395, 1999.
- [25] M. L. Fredman. How good is the information theory bound in sorting? *Theor. Comput. Sci*, 1(4):355–361, 1976.
- [26] M. L. Fredman. New bounds on the complexity of the shortest path problem. *SIAM J. Comput.*, 5(1):83–89, 1976.
- [27] A. Freund. Improved subquadratic 3SUM. *Algorithmica*. First Online, October 19, 2015.
- [28] A. Gajentaan and M. H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom.*, 5:165–185, 1995.

- [29] A. Grønlund and S. Pettie. Threesomes, degenerates, and love triangles. In *Proceedings of the 55th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 621–630, 2014.
- [30] J. Hartmanis and R. E. Stearns. On the computational complexity of algorithms. *Trans. Amer. Math. Soc.*, 117:285–306, 1965.
- [31] A. Hernández-Barrera. Finding an $o(n^2 \log n)$ algorithm is sometimes hard. In *Proceedings of the 8th Canadian Conference on Computational Geometry*, pages 289–294, 1996.
- [32] R. Impagliazzo and R. Paturi. On the complexity of k -SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, Mar. 2001.
- [33] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.
- [34] Z. Jafargholi and E. Viola. 3SUM, 3XOR, triangles. *CoRR*, abs/1305.3827, 2013.
- [35] T. Kopelowitz, S. Pettie, and E. Porat. 3SUM hardness in (dynamic) data structures. *CoRR*, abs/1407.6756, 2014.
- [36] F. P. Preparata and M. I. Shamos. *Computational Geometry*. Springer, New York, NY, 1985.
- [37] M. Pătraşcu. Towards polynomial lower bounds for dynamic problems. In *Proceedings 42nd ACM Symposium on Theory of Computing (STOC)*, pages 603–610, 2010.
- [38] M. Pătraşcu and R. Williams. On the possibility of faster SAT algorithms. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1065–1075, 2010.
- [39] L. Roditty and V. V. Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, pages 515–524, 2013.
- [40] M. A. Soss, J. Erickson, and M. H. Overmars. Preprocessing chains for fast dihedral rotations is hard or even impossible. *Comput. Geom.*, 26(3):235–246, 2003.
- [41] O. Weimann, A. Abboud, and V. V. Williams. Consequences of faster sequence alignment. In *Proceedings of the 41st Int’l Colloquium on Automata, Languages, and Programming (ICALP)*, pages 39–51, 2014.
- [42] V. V. Williams and R. Williams. Subcubic equivalences between path, matrix and triangle problems. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 645–654, 2010.
- [43] V. V. Williams and R. Williams. Finding, minimizing, and counting weighted subgraphs. *SIAM J. Comput.*, 42(3):831–854, 2013.